Amrutha Kanakatte Ravishankar

201-705-6654 amruthakravishankar@outlook.com linkedin.com/in/amrutha-kanakatte-ravishankar-/
github.com/harithsya24 harithsya24.github.io/profile/

Education

Stevens Institute of Technology **Jan. 2024 – Dec. 2025**
Master of Science in Computer Science, GPA: 3.88 *Hoboken, New Jersey*
- Relevant Coursework: Machine Learning, Deep Learning, Natural Language Processing, Generative AI, Knowledge Discovery and Data Mining, Statistical Machine Learning

JSS Science and Technology University **Aug. 2019 – Aug. 2023**
Bachelor of Engineering in Electronics and Communication Engineering, GPA: 8.94 / 10 *Mysuru, India*
- Relevant Coursework: Digital Signal Processing, Linear Algebra, Computer Networks

Experience

PONS **June 2025 – Present**
AI Research Engineer Intern *Newark, NJ*
- Developed GAN-based ultrasound image enhancement system with SPADE normalization in PyTorch, improving diagnostic clarity by 25% and processing 15,000+ images/hour
- Trained on 70% of image dataset; deployed scalable ML inference pipeline using Docker, Kubernetes, and GCP (Compute Engine, Vertex AI, Cloud Storage)
- Optimized U-Net architecture for tissue segmentation, achieving 87% accuracy in medical image boundary detection

Stevens School of Business **May 2025 – Present**
Graduate Research Assistant *Hoboken, NJ*
- Designed FinVoc2Vec, a financial audio-to-vector model using Wav2Vec 2.0 with transformer and convolutional layers, achieving 85% sentiment classification accuracy (up from 54%)
- Built full-stack system using Python and Flask to convert financial voice notes into structured intelligence for compliance and investment insights
- Automated ingestion of 500+ financial documents daily via Google Gemini API, reducing manual effort by 60%

FutureWave **Aug. 2021 – May 2022**
Software Engineer Intern *Mysuru, India*
- Reduced crash rate by 20% using automated testing (PyTest) and robust error logging in Flask-based systems
- Improved release quality by reducing production bugs 30% through enhanced code reviews and integration test coverage

Projects

CloudMart: Multi-Cloud E-commerce Platform | *AWS, GCP, Azure, Terraform, GitHub* **January 2025**
- Architected fault-tolerant multi-cloud platform using AWS (Lambda, Bedrock, EKS, DynamoDB), GCP (BigQuery), Azure (AI Language), supporting 10,000+ concurrent users
- Created CI/CD pipeline with AWS CodePipeline, ECR, and Terraform—reduced deployment time from 2 hours to 25 minutes, eliminating 95% of manual errors

Agentic Medical RAG System for Clinical Decision Support | *RAG, FAISS, PubMed API* **December 2024**
- Built production-ready RAG system indexing 1.2M PubMed articles and Hugging Face datasets using FAISS; achieved 87% query relevance and 0.3s response time
- Deployed Flask-based interface for diagnostic, treatment, and drug-related queries with 95% retrieval precision

Technical Skills

Programming: Python, C++, C, SQL
ML & AI Frameworks: PyTorch, TensorFlow, Scikit-learn, Transformers, Hugging Face, Sentence Transformers
AI Specialties: RAG Systems, Diffusion Models, Computer Vision, Multi-modal AI, Prompt Engineering, Model Optimization, Knowledge Distillation
Cloud & DevOps: AWS (SageMaker, Lambda, EC2, S3), GCP (Vertex AI, BigQuery), Docker, Kubernetes, CI/CD, Terraform
Databases & Tools: MongoDB, FAISS, Redis, Pinecone, FastAPI, NumPy, Pandas, OpenCV, Git

Publications

Comparative Analysis of Transformer Architectures for Intrusion Detection Systems **August 2024**
IEEE NMITCON 2024
Deep Learning Approaches for Audio Source Separation in Karaoke Applications **June 2023**
IJIRAE, Vol. 10, Issue 06